Point of View

THE CYBERSECURITY SPEED GAP | INDUSTRY SIGNAL

WHEN ATTACKING GETS CHEAPER THAN DEFENDING

AI is moving parts of vulnerability discovery, reconnaissance and attack experimentation towards machine speed. Enterprise security still runs on asset inventories, owners, testing windows and approval queues. The gap between those two operating speeds is becoming the new asymmetry.



The most telling cyber story today is not a single breach. It is the pattern: AI-assisted research uncovering long-standing flaws in widely used software; agents crossing the boundaries of an authorised test; credentials exposed through an AI software supply chain; and conventional weaknesses being found and revisited at a pace that human teams struggle to match. The failure modes differ. The imbalance behind them is increasingly similar.

Six Signals
and What They Actually Establish

WHERE / WHEN	WHAT HAS BEEN REPORTED	WHAT THE EVIDENCE ESTABLISHES
Global · Oct 2026	1. AI helps defenders find flaws faster. OpenAI/ Wiz identified two critical PHP flaws; Amazon saw an exploitation attempt on 17 October, before disclosure. [1,2]	AI assisted, defender-led discovery is established; speed of AI assisted discovery is established.
India · Oct 2026	2. Agents can exceed task authority. CERT-In observed research agents reach live systems without permission. [4]	The activity was disclosed by model providers; affected organizations did not detect the unauthorized activity themselves.
India · Oct 2026	3. Agentic tooling compresses attack paths. CERT-In observed financially motivated intrusions using open-source agents, with access-to-impact reduced to hours. [4]	Agentic tooling reduced access-to-impact time to hours; the intrusions involved exposed keys and dormant APIs. permissive roles were also present.
Global · Oct 2026	4. AI gateways are supply-chain targets. A compromised LiteLLM v1.107.0 release reached public package/container registries on 21 October. [7]	While the evidence does not establish compromise of every LiteLLM user, the affected artifacts contained malware with credential theft and remote-code execution capabilities.
Global · 2026	5. Agent adoption adds delegated risk. CISA partners warn of privilege creep, expanded attack surfaces and limited visibility into agent actions. [6]	These are identity, architecture and operating-model issues, not just prompt-safety issues.
Global · 2026	6. Conventional flaws remain central. Verizon: 31% of breach entries used exploits; a separate AI-use dataset found a median of 15 ATT&CK techniques researched per actor. [8]	Conventional flaws remain a material breach route. The cited AI-use data shows research activity, not confirmed AI-driven exploitation.

> A note on the evidence: discoveries, exploitation attempts, production intrusions, vendor incidents and reported financial losses are different categories of evidence. They should not be collapsed into a single “AI breach” narrative. They should also not be treated as unrelated. Taken together, they show that the window between a weakness becoming usable and an organization proving it is safe is getting harder to manage.

The question to ask: if an agent can discover, test and iterate in hours, why does our defense still move in weeks?

AI Changes the Economics of Repeated Attacks

Cybercrime has always been asymmetric. The defender must protect a complex estate, every day, across an expanding attack surface. The attacker needs one workable path, once. AI improves that trade-off at the margin for repeatable work. Operators can use models to sift code, test variants, draft tailored lures, interpret failed attempts and delegate routine work to automated workflows. The same capabilities can be reused across many targets.

That does not make every attack autonomous or every exploit cheap. Skilled operators, infrastructure and access still matter. The practical shift is that each reconnaissance cycle, phishing variant or exploit experiment may require less specialist time, even where access and exploitation still demand expertise. Defense remains enterprise-specific. Teams still need to know which asset is exposed, what business process depends on it, what a change could disrupt and who has authority to act.

Where the Shift Shows Up: Discovery, Execution and Response

Discovery. The PHP2Shell episode is instructive. OpenAI and Wiz reported AI-assisted defender-led discovery of critical PHP flaws and coordinated remediation. Separately, Amazon observed an exploitation attempt against a honeypot on 17 October, four days before disclosure. No evidence ties that attempt to AI-generated exploitation or proves a production compromise. The defensive clock can start before a CVE is publicly disclosed [1,2]

Execution. Google Threat Intelligence observed agent-enabled credential harvesting planned and executed in under six hours after a cloud compromise. Separately, CERT-In reported financially motivated intrusions in India where agentic tools compressed the time from initial access to impact from days to hours. These specific cases do not prove that every attacker operates autonomously. [3,4]

Response. Buying more detection alone is insufficient. A vulnerability can be discovered, triaged, assigned, tested and declared closed even while the attack path remains open. Effective organizations link discovery to decision authority, remediation and independent verification rather than keeping each step in a separate queue.

The Exposure-to-Verified-Fix Interval

Time to detect an incident is only one measure. The full interval runs from a weakness becoming discoverable, to the organization knowing whether it is exposed, containing the route, deploying a durable fix and proving that the path is closed. Every hand-off is time an adversary may not need.

AI-assisted defense carries its own risk. The same models that help find more weaknesses can overwhelm a team that has not agreed how to prioritize them or who can authorize the response. A larger vulnerability backlog is not a safer enterprise. The gap narrows only when discovery and remediation run as one process—with one accountable owner, one clock and a retest at the end.

> Google observed an agent-enabled attack operation in under six hours; Verizon reports a 43-day median remediation time in its studied vulnerability population. These are different datasets and not a direct comparison, but they illustrate why enterprises must shorten the path from warning to verified closure. [3,8]



Security Operating Models Still Run on Human Time

Many security operating models still rely on periodic scans, fragmented ownership, human- scale investigation and change windows measured in days. These handoffs struggle with agents that can examine an exposed estate continuously, try alternate paths, switch tools and retry without fatigue. The challenge is not to replace proven frameworks. It is to speed up accountable decisions from first signal at 2 a.m.

Five Fault Lines in Enterprise Defense

1 **The unknown estate.** Critical exposure is not limited to the applications that teams know they own. It may sit in a legacy portal, a vendor-operated microsite, an inherited application, a test API or an administrative endpoint. An inventory that cannot connect an asset, its owner, identity and business consequence is not an operational control.

2 **The time lag.** A vulnerability publication is not necessarily the first time an attacker tried it. Emergency changes still wait on testing, providers, approvals and change boards. “Patched” should mean deployed, restarted where necessary and independently verified, not merely that a ticket was closed.

3 **The delegated-authority gap.** Agents can hold legitimate access and still act outside the intended task. If authorization exists only in a prompt, it is not a reliable control. Every agent needs a short-lived identity, task-scoped permissions, default-deny egress, downstream authorization and a clear stop condition.

4 **The dependency problem.** A trusted package, model, plugin, certificate or credential can become a distribution route. The LiteLLM incident is a reminder that signing, approvals and release gates are only as strong as the identities and logs behind them. [7]

5 **The adversary-intelligence gap.** Security teams can miss new AI-enabled attack patterns: parallel agents, tool chaining and persistent retries. Use AI-aware threat intelligence and adversary emulation to test responses.

Software That Can Act

Agent security is more than prompt safety. OWASP describes excessive agency as a combination of excessive functionality, excessive permissions and excessive autonomy. CISA and international partners have also warned about privilege creep, expanded attack surfaces and weak event records in agentic systems. A prompt filter is not a substitute for network controls, scoped identities, output validation and an independent approval step for high-impact actions. The model should be treated as an untrusted decision component—not as the policy enforcement point. [5,6]

Enterprises now face two pressures at once. Adversaries use AI to probe traditional weaknesses. Employees and applications are also giving agents access to code, records and business workflows. An agent that can read email, query a data warehouse, modify code or call a payment API can turn a data-access problem into an action problem. The security perimeter now extends from people and applications to delegated actors. Identity must answer not only who is calling, but what the agent may do now, for which task, with what data and with what approval.

Why This Belongs on the Board Agenda

The Indian incidents described by CERT-In are a useful warning because they do not depend on exotic vulnerabilities. They involved exposed credentials, forgotten APIs, permissive cloud roles and internet-facing management surfaces—the same weaknesses many organizations already know they need to fix. That is precisely the point: AI does not need to invent a new weakness if it can find and reuse an old one faster. For Indian businesses, cyber resilience, DPDP safeguards and third-party accountability should be discussed together whenever an agent can access personal or confidential data. [4]

> A larger security budget does not automatically create a faster security organization. The fault lines sit in asset reality, decision rights, delegated authority and trusted dependencies—areas that technology can expose, but leadership must resolve.

Shorten the Time from First Signal to Verified Fix

Buying an “AI SOC” alone will not solve this asymmetry. Leaders need evidence of a single accountable exposure-to-closure process: discover, determine impact, authorize containment, fix and independently verify. Within 90 days, pilot the end-to-end process across ten critical external services and suppliers. Assign decision rights to the CISO, business, application and infrastructure owners. Start the clock when the first actionable warning appears and not when a CVE is published.

Four Priorities Before Buying Another Tool

1. **Make exposure visible.** Map every critical external service to its business owner, technical owner, provider and business consequence, including test systems and suppliers. Continuously monitor for exploitation attempts, unexpected processes, new files and unusual outbound connections. Route urgent alerts to a named, empowered responder.

2. **Measure time to safety.** Track five clocks: time to detect an exploitation attempt, identify an exposed asset, contain the route, deploy a durable fix and verify it. Prioritize by reachability and consequence, not CVSS alone. Pre-approve interim containment and rehearse emergency change authority. Publish the five clocks to the executive risk dashboard.

3. **Link every finding to a fix.** Use AI-assisted testing, code analysis and SOC investigation to match the adversary’s tempo, but couple every finding to an accountable remediation workflow, a decision deadline and a retest. Hold one owner accountable for closure; report attack paths verified as closed, not just findings raised.

4. **Govern what agents do.** Maintain an inventory of production agents, tools, data stores and non-human identities. Restrict permissions and egress; treat retrieved content as untrusted; require independent approval for irreversible actions; log actions outside the model; monitor behavior against a baseline; rehearse a kill switch; and test agent-aware adversary scenarios.

Five Questions for the Next Risk Committee

Which are our ten most consequential internet-facing services, and can we show who owns each one? How quickly can we determine whether a newly disclosed weakness affects them? Who is alerted at 2 a.m., and who can authorize containment? Can we prove that an agent cannot move from a read task to a write or send action without an independent decision? Can we prove that a fix closed the attack path rather than only closing the ticket?

> AI has not made conventional cybersecurity obsolete. It has made the weaknesses we used to tolerate more dangerous. The marginal effort to repeat parts of an attack may fall. The enterprise coordination required for a defensible response remains. The strongest organizations will not be those with the newest tools. They will be the ones that shorten the distance between a warning, an authorized decision and a verified fix.

Selected Evidence and Source Status: revised 4 November 2026

[1] OpenAI — Coordinated disclosure of critical vulnerabilities

[2] Amazon Threat Intelligence — PHP2Shell exploitation prior to disclosure

[3] Google Threat Intelligence — From prompting to autonomy

[4] CERT-In Advisory CIAD-2026-0049 — AI Agents as Attack Vectors

[5] OWASP — LLM06:2025 Excessive Agency

[6] CISA and international partners — Careful Adoption of Agentic AI Services

[7] LiteLLM — 23 Oct canonical investigation of compromised v1.107.0 release

[8] Verizon — 2026 Data Breach Investigations Report

About Uniqus Consultech:

Uniqus Consultech is an AI and global tech-enabled consulting company that specializes in Accounting & Reporting Consulting, Governance, Risk & Compliance, Sustainability & Climate Consulting, Tech Consulting, and Valuations. The Company is co-founded by consulting veterans Jamil Khatri and Sandip Khetan and backed by marquee investors such as Nexus Venture Partners, Sorin Investments, and UST.

Uniqus has a global team of 800+ professionals, led by 100+ Partners and Directors, across 13 offices in the USA, the Middle East, and India. The company serves more than 400 clients, including marquee names in each of the markets it operates in.

Uniqus is committed to leveraging technology and an integrated global delivery model to provide best-in-class consulting services to its clients.

About Our Tech Consulting Services:

Technology is now a core driver of competitiveness, resilience, and enterprise value.

As organizations digitize operations end-to-end, technology environments are becoming more complex. Fragmented systems, inconsistent data, evolving regulatory expectations, cyber threats, and rapid AI adoption can slow transformation and increase enterprise risk.

Our Tech Consulting practice helps clients cut through this complexity with clarity, discipline, and execution focus. We combine deep technology capabilities with differentiated functional expertise across finance, risk, governance, sustainability, and compliance. Our services help organizations modernize applications, strengthen cyber and technology risk management, build trusted data foundations, automate processes, and adopt AI optimally.

We believe durable transformation requires more than technology implementation. It requires secure architecture, governed data, embedded controls, human oversight, and measurable business outcomes. By integrating these principles from day one, we help clients innovate with confidence and scale with control.

Scan this code to find
more content like this

Visit us at *www.uniqus.com*

or follow us on